

12-13-99

A

12/10/99

THOMASON, MOSER & PATTERSON

ATTORNEYS AT LAW

THE GALLERIA
2 BRIDGE AVENUE
POST OFFICE BOX 8160
RED BANK, NEW JERSEY 07701
TELEPHONE (732) 530-9404
TELEFAX (732) 530-9808

PATENT APPLICATION

Assistant Commissioner for Patents
Box **Patent Application**
Washington, D. C. 20231

Sir:

Enclosed herewith for filing is the following **utility patent application which claims priority to U.S. Provisional Application No. 60/127,337, filed April 1, 1999:**

Applicant(s): Tom Giammarresi

Title of application: METHOD AND APPARATUS OF LOAD SHARING AND
FAULT TOLERANCE IN AN INTERACTIVE VIDEO DISTRIBUTION SYTSTEM

Pages of specification: 21 (14 + 6 pages of claims and 1 page
abstract)

Sheets of drawing: 7

Executed on: 12/6/99

Attorney Docket No. 533/047

PATENT APPLICATION FILING FEE CALCULATION

	<u>No. Filed</u>	<u>Less</u>	<u>Rate/Claim</u>	<u>Fee</u>
Total Claims	21	-20	1 x \$ 18.00	\$ 18.00
Independent Claims	2	- 3	0 x \$ 78.00	\$ 0.00
Minimum Filing Fee				\$ 760.00
Multiple Dependency Fee (if applicable - \$270.00)				0.00
TOTAL FILING FEE AT LARGE ENTITY RATE				\$ 778.00
50% Reduction for Small Entity (Independent Inventor, Non-profit Corporation, or Small Business Concern) - appropriate verified statement attached				\$ 389.00
TOTAL FILING FEE AT SMALL ENTITY RATE				\$ 389.00

Kindly charge the entire and/or any additional filing fee, as appropriate, to deposit account number 20-0782. To facilitate that charge, a duplicate copy of this letter is enclosed herewith.

Also enclosed herewith for filing in connection with the enclosed application are:

☐ Oath;
☒ Declaration and Power of Attorney;
☐ Disclosure Statement;
☐ Letter referencing previously filed disclosure document; number _____ filed _____;
☒ Verified Statement claiming small entity status;
☒ Assignment of the application to: DIVA Systems Corporation
☒ Claim(s) to priority
 Serial Number Filing date
 60/127,337 April 1, 1999
☐ A certified copy of a _____ patent application or inventor's certificate, filed _____ and serial no. _____, upon which a claim to priority is made;

Respectfully submitted,

12/10/99

E. J. Wall

Eamon J. Wall, Attorney
Reg. No. 39,414

Thomason, Moser & Patterson, LLP
Attorneys At Law
2-40 Bridge Avenue
P.O. Box 8160
Red Bank, NJ 07701

*****EXPRESS MAIL CERTIFICATION*****

"Express Mail" mailing label number: EL530945351US
Date of deposit: 12/10/99

I hereby certify that this patent application and related papers is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner of Patents, Box Patent Application, Washington, D.C. 20231.

Laura E. Crater
Signature of person mailing paper or fee

LAURA E CRATER
Name of person mailing paper or fee

**VERIFIED STATEMENT (DECLARATION) CLAIMING SMALL ENTITY
STATUS (37 CFR 1.9(f) AND 1.27 (c)) - SMALL BUSINESS CONCERN**

Docket No.
533/047

Serial No.

Filing Date

Patent No.

Issue Date

Applicant/

Patentee: **Tom Giammarresi**

Invention:

METHOD AND APPARATUS OF LOAD SHARING AND FAULT TOLERANCE IN AN INTERACTIVE, etc.

I hereby declare that I am:

- ☐ the owner of the small business concern identified below:
☒ an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF CONCERN: _____

ADDRESS OF CONCERN: _____

I hereby declare that the above-identified small business concern qualifies as a small business concern as defined in 13 CFR 121.3-18, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees under Section 41(a) and (b) of Title 35, United States Code, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the above identified invention described in:

- ☒ the specification filed herewith with title as listed above.
☐ the application identified above.
☐ the patent identified above.

If the rights held by the above-identified small business concern are not exclusive, each individual, concern or organization having rights to the invention is listed on the next page and no rights to the invention are held by any person, other than the inventor, who could not qualify as an independent inventor under 37 CFR 1.9(c) or by any concern which would not qualify as a small business concern under 37 CFR 1.9(d) or a nonprofit organization under 37 CFR 1.9(e).

Each person, concern or organization to which I have assigned, granted, conveyed, or licensed or am under an obligation under contract or law to assign, grant, convey, or license any rights in the invention is listed below:

- ☒ no such person, concern or organization exists.
☐ each such person, concern or organization is listed below.

FULL NAME _____

ADDRESS _____

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

FULL NAME _____

ADDRESS _____

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

FULL NAME _____

ADDRESS _____

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

FULL NAME _____

ADDRESS _____

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING:

Stephanie A. Storms

TITLE OF PERSON SIGNING _____

OTHER THAN OWNER:

Senior Vice President/General Counsel

ADDRESS OF PERSON SIGNING:

c/o DIVA Systems Corporation800 Saginaw DriveRedwood City, CA 94063

SIGNATURE:

Stephanie A. Storms

DATE:

12/5/99

METHOD AND APPARATUS OF LOAD SHARING AND FAULT TOLERANCE IN AN INTERACTIVE VIDEO DISTRIBUTION SYSTEM

CROSS REFERENCE TO RELATED APPLICATIONS

- 5 This application claims the benefit of U.S. Provisional Application No. 60/127,337, filed April 1, 1999, which is hereby incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

1. Field of Invention

- 10 The present invention relates to an interactive information distribution system. More particularly, the invention relates to a method and apparatus for improving sharing of the processing loads and increasing the fault tolerance at each of a plurality of head-ends in an information distribution system.

2. Description of the Background Art

- 15 Video distribution systems typically utilize a plurality of cable head-ends. A head-end serves as a distribution point for a designated "neighborhood" of subscribers. Subscriber requests for video information, such as movies, are made by a subscriber using a remote control device to select from a menu of available titles displayed on a display device. After selection by a subscriber, a request for the selected video
20 information is sent to the local head-end supporting the subscriber. The requested video information is then transmitted from the head-end to the subscriber for viewing. A typical head-end comprises a video server system that contains subsystems for managing, storing and distributing the video content.

- 25 In a system comprising a large number of subscribers, each head-end experiences a high level of activity. This typically results in delays in responding to subscriber requests. In addition, data processing and/or transmission errors increase as the demand placed upon the head-end increases. For example, physical hardware errors, such as those caused by thermal stresses created during component over-utilization, may result in unacceptable viewing experiences for the subscriber.

- 30 Therefore, there is a need in the art for more robust head-end functionality within an information distribution system by distributing head-end processing functions among a

plurality of processing elements in a manner that provides increased fault tolerance and, ideally, reduced system latency.

SUMMARY OF INVENTION

5 The disadvantages heretofore associated with the prior art are overcome by the present invention of an apparatus and method suitable for use in an information server, such as a video-on-demand system head-end. The apparatus and method provide processor load sharing by incorporating redundant hardware and software at the head-end. In this manner, redundant hardware and software at the head-end may share processing duties, as well as
10 advantageously increase fault tolerance, so as to minimize the negative effects of a failure at one of the redundant components.

In an exemplary embodiment of the invention, a system head-end comprises subsystems that are used to manage, store, and distribute video content. These subsystems include a stream server, a video content storage device, numerous managing modules
15 providing various system management functions, and at least two head-end controllers to process and store session-state data executed by the managing modules. The data stored at each head-end controller is available for access by the other head-end controllers by either replicating the data onto a storage device of the other head-end controller or linking each storage device through a network to provide shared data access.

20 The data processing and delivery load is divided evenly amongst the head-end controllers. Each head-end controller is dedicated to a plurality of access controllers. The access controllers are in turn dedicated to a plurality of subscriber set-top boxes. The set top boxes provide a subscriber access to the VOD system as well as to decode video information supplied to the subscriber. In the preferred embodiment having two head-end
25 controllers, a first head-end controller provides session-state management for one-half of the access controllers and their respective subscriber set-top boxes, while the second head-end controller provides session-state management for the remaining access controllers and their respective set-top boxes.

In the event of a failure in either head-end controller, the remaining operational
30 head-end controller will service all of the access controllers and corresponding set-top

boxes. Successful operations are accomplished by the processing of the data from the failed head-end controller, which was either mirrored onto the operational head-end controller's storage device prior to failure, or stored on a shared storage device through a network. Thus, the method and apparatus advantageously increases the fault tolerance at the head-end by sharing the processing loads amongst head-end controllers, as well as eliminating a single fault failure at the hardware level of the head-end controllers through redundancy.

BRIEF DESCRIPTION OF THE DRAWINGS

- 10 The teachings of the present invention can be readily understood by considering the following detailed description in conjunction with the accompanying drawings, in which:
- FIG. 1A and FIG. 1B together depict a high-level block diagram of a first embodiment of an interactive information distribution system;
- FIG. 2A and FIG. 2B together depict a high-level block diagram of a second
15 embodiment of an interactive information distribution system;
- FIG. 3A and FIG. 3B together depict a flowchart of a method that facilitates fault tolerance at a head-end of an interactive information distribution system; and
- FIG. 4 depicts a flowchart of a method for continuing a subscriber session during a failure mode of operation at a head-end controller.
- 20 To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures.

DETAILED DESCRIPTION OF THE INVENTION

FIG. 1 depicts a high-level block diagram of an interactive information distribution
25 system. FIG. 1 is formed by arranging FIG. 1A and FIG. 1B according to the graphics depicted in the respective figures.

The interactive information distribution system 100 of FIG. 1 comprises a head-end 101, a transport subsystem or network 110 and subscriber equipment 124. The head-end 101 receives subscriber requests for information such as movies or other content, and

responsively provides or "streams" the content to the appropriate subscriber equipment 124 via a network 110.

A plurality of subscriber equipment 124-1 through 124-x (collectively subscriber equipment 124) each comprise a set-top box (STB) 118, an input device (e.g., remote control) 120 and a display device 122. Communications between the head-end 101 and the subscriber equipment 124 are transmitted across a transport system network 110 by either cable or telephone transport mediums.

The transport network 110 is typically, but not exclusively, a conventional bi-directional, hybrid fiber-coax cable network. Depending upon the fiber node size, one embodiment of the invention utilizes two to five conventional cable channels (e.g., 6 MHz bandwidth channels). Each channel is capable of downstreaming, via a plurality of downstream information channels 108_1 through 108_n (collectively information channels 108), up to 10 streams of video information per channel at the same time. Assuming a 10 to 1 concentration, i.e., not all subscribers are viewing at the same time, then approximately 500 potential subscribers may be connected to each node. In addition to downstream information channels 108, the transport network 110 must also support downstream command channels 107_1 through 107_n (collectively command channels 107) and upstream "back" channels 109_1 through 109_n (collectively back channels 109).

The head-end 101 provides control information for transmission through the downstream command channel 107 contained within the cable transport network 110 to the subscriber equipment 124. This command and control information is transmitted illustratively, on a carrier in the range 50 to 750 MHz using a 1 MHz bandwidth, i.e., the command information is frequency multiplexed with the information channel and transmitted through the network 110. Furthermore, the subscriber equipment 124 communicates via the reverse (or back) channel 109 to the head-end 101 through the cable transport network 110. Each reverse channel carries, for example, a BPSK modulated signal on a carrier in the range 5-42 MHz, where the channel capacity is approximately 64 Kbps. Other modulation formats (e.g., QPSK at 1.5 Mbs) or channel capacities may be used. Alternatively, the plurality of data streams is multiplexed onto an optical fiber (a trunk) and each head-end 101 is connected to the trunk by a "drop line."

The subscriber equipment 124 is coupled to the transport network 110 by signal path (e.g., coaxial cable) 105-1 through 105-x (collectively signal path 105) via the set top box 118. The set top boxes 118 receive and demodulate the downstream signals including those propagated through both the command channel 107 and the information channel 108.

5 The set top boxes optionally demodulate standard cable television signals received from the network. Thus, a single set top box can be used to receive all the cable services provided by the network. The set top boxes also provide interactive control of the information presentation. The presentation is controlled via the input device 120, e.g., an infrared (IR), radio frequency (RF), or some other remote control unit. The information, e.g., audio,

10 video, still photographs, graphics, and other multimedia programs and the like are portrayed on the display device 122 such as a television, video monitor, stereo system, or otherwise.

The viewer selects entry into the system by manipulating the buttons (or joystick) of the remote control device 120. During the selection process, the viewer (now a potential

15 subscriber) is presented with an on-screen browser (a graphical user interface) that aids the viewer in finding information, the prices of the selections, search aides, and the like. The commands used to navigate throughout the various menus are transmitted from the set top terminal 118 through the signal path 105 to the head-end 101 via the back channel 109. The head-end 101 responds to customer commands through the signal path 105 via the

20 downstream command 107 and information channels 108. When the subscriber executes the navigator, a "session" is open for that particular subscriber.

The head-end 101 of the interactive information distribution system 100 comprises a stream server 102, a content storage device 106, at least two head-end controllers 130-1 and 130-2 (collectively head-end controllers 130), and a plurality of access controllers (AC)

25 140-1 through 140-m (collectively access controllers 140). The content storage device 106 is coupled to the stream server 102, thereby providing a repository of video assets that are available for request by the subscribers. The content storage device 106 typically contains thousands of video assets from which the subscriber may choose.

The video stream server 102 is coupled to each head-end controller 130 through

30 signal paths 111 to allow communications of command information between the server 102

and each head-end controller 130. Additionally, the video stream server 102 is coupled to each access controller 140 to provide a plurality of packetized data streams via a signal path 104 and to provide a synchronization clock signal via signal path 103. The packetized data streams contain isochronous information as well as movies or other video assets
5 retrieved from the video content storage device 106.

The plurality of access controllers 140 are coupled to each of the head-end controllers 130. This coupling 112 may illustratively be Ethernet or fiber channel cabling. Additionally, the plurality of access controllers 140 are coupled to the cable transport subsystem 110 through the downstream command channel 107, the information channel
10 108 and the back channel 109.

Furthermore, each of the plurality of access controllers 140 comprises multiplex boards, modems and other apparatus (not shown). The access controllers are primarily responsible for controlling the quadrature amplitude modulation (QAM) modulators and multiplexing boards that are used as mediums for transmitting and receiving the data
15 between the subscriber equipment 124 and the head-end controllers 130.

Specifically, the access controllers 140 control the provisioning of video information between the server 102 and subscriber equipment 124. In the event of a head-end controller failure, the access controllers control the switching over to the operative head-end controller. Thus, by including an additional head-end controller at each head-end,
20 fault tolerance for the system is improved.

The number of head-end controllers 130 is proportional to the number of subscribers being serviced by the system. Each head-end controller can generally service up to 500 subscribers. Additional head-end controllers 130 may be added to a head-end 101 as required.

Each of the head-end controllers 130 of the head-end 101 comprises various managing modules 132-1 and 132-2 (collectively managing modules 132), a processor 135-1 and 135-2 (collectively processors 135), and memory (e.g., RAM) 136-1 and 136-2 (collectively memory 136). A Graphic User Interface (GUI) workstation 134-1 and 134-2 (collectively GUI workstations 134) is also available for operator interaction. Additionally,
25 FIG 1A depicts a plurality of local storage devices 137-1 and 137-2 (collectively local
30

storage device 137) also coupled with each head-end controller 130. Alternatively, FIG. 2A depicts a second embodiment having a central storage device 139 networked to each head-end controller 130.

The processor's 135 function is to process session-state data that is executed by the
5 managing modules 132. The optional GUI workstation 134 is a computer interface to allow interaction with an operator.

The processor 135, memory 136 and various managing modules 132 of each head-end controller 130 are coupled together on a circuit board such as, a single compact PCI board, and stored in a rack unit. Thus, the rack unit will contain at least two head-end
10 controllers 130 with the ability to hold additional boards as required.

The storage devices 137 and 139 depicted in FIGS. 1A and 2A may illustratively be either fiber channel or SCSI hard drives. Each local storage device 137, as illustrated in FIG. 1A, is coupled by a network connection to each head-end controller 130. Thus, each head-end controller 130 may access more than one storage device 137, if required.

15 The networked storage device 139, as illustrated in FIG. 2A, is a central storage device at the head-end 101, and connected by a network to each head-end controller 130. Thus, each head-end controller 130 shares a single storage device 139. By either method, each head-end controller 130 has access to the data stored on the storage device by the other head-end controller 130.

20 The various managing modules 132 within each head-end controller 130 are programs that execute session-state data. A session is the interactivity between the subscriber using their subscriber equipment 124 and the head-end 101, whereby the processing of a video request is provided. Session-state data is information that defines the state of the session. It includes who the subscribers are, which set-top boxes are active,
25 what video asset is being watched, the addresses of the set-top boxes, which modulators are being used, which navigation screen the subscriber is watching, and similar session related information.

Data that is non-volatile is deemed as permanent session-state data that may be required frequently, long term, or in the event of a power loss. Such permanent data is
30 written to a local storage device 137 of FIG. 1A, or the centrally networked storage device

139 of FIG. 2A. That is, once processed by the processor 135, the managing module stores the session-state data on its respective local storage device 137 or the centrally networked storage device 139.

During operation of the preferred embodiment of a head-end 101 having two head-
5 end controllers 130-1 and 130-2, session-state data is permanently stored by writing the data to the local disc drive 137-1 (FIG. 1) that is assigned to head-end controller 130-1, and then replicating the data image periodically to the other local storage device 137-2 on secondary head-end controller 130-2. In effect, each head-end controller's local storage device has a mirror image of the other local storage disk. In the case where more than two
10 head-end controllers 130 exist at a head-end 101, the local storage device 137 for each head end controller may be a shared device centrally located at the head-end 101.

Writing of the session-state data to the local storage devices 137 is performed via the managing module that is currently managing the data. Although writing of the session-state data may be performed in real time, it is not necessary since each storage device 137 is
15 periodically updated with the session-state data from the other storage device. Moreover, the periodic updates occur in frequent intervals, thus allowing enough session-state data from the primary head-end controller 130 to be stored on the storage device 137 of the secondary head-end controller to avoid a system "crash" in the event of a primary head-end controller 130 failure.

20 However, the session-state data must be accessed in real time. After the session-state data is stored, access to the fixed storage devices may be made through the managing module requiring the information. By utilizing a local storage device 137 for each head-end controller 130, only a minimal amount of session-state data will be lost in the event of a primary head-end controller failure.

25 Alternatively, as shown in FIG. 2, the system 100 permits the processor 135 to directly access the centrally networked storage device 139 to read the required session-state data. This latter method is much faster and thereby assists in keeping the processing of data synchronized and accessed in real time. Furthermore, by utilizing a centrally networked storage device 139, no session-state data will be lost in the event of a primary head-end

failure, since the system does not have to rely on frequently made content updates to other storage devices.

Volatile session-state data is data that does not require permanent storage on a permanent storage device 137. Rather, after execution by a managing module, the session-state data is stored temporarily in the RAM 136 by the processor 135 and later discarded. Referring to FIG. 1, non-permanent session-state data executed by the managing modules 132-1 at the primary head-end controller 130-1 is stored on the corresponding memory 135-1. Copies of the non-permanent session-state data are made and transferred periodically in frequent intervals during each session instance, from the current processing head-end controller (i.e., primary head-end controller 130-1) to the other head-end controller (i.e., secondary head-end controller 130-2) as a method of redundant data storage. In an instance where the primary head-end controller 130-1 fails, the secondary head-end controller 130-2 will utilize the most current session-state data from the primary head-end controller 130-1. Updating the RAM at each head-end controller 130 is frequent enough to avoid a system crash in the event of such primary head-end controller failure.

Synchronization at the memory may be accomplished by writing and reading to and from the RAM 136 directly from the processor 135 without having to go through the managing module 132. The data in the RAM 136 at each head-end controller 130 may be transferred illustratively, via the Ethernet transport medium 112 coupling the primary and secondary head-end controllers 130, or via a common bus between the processors 135 and RAM 136 of both head-end controllers. With each update, the secondary head-end controller 130-2 stores this information on the corresponding RAM 136-2, thereby having the executed session-state data available should a fault occur in the primary head-end controller. Thus, each memory device 136 is updated with a mirror image of the memory devices from other head-end controllers 130.

Referring back to FIG. 1A, each head-end controller 130 has a variety of managing modules 132 stored thereon. Each managing module has a distinct function for managing and processing specific data at different times. For example, a portion of the managing modules are dedicated to processing session-state data, that is generated during the subscriber's requests for video content. Other managing modules manage video asset

allocation and storage at either the head-end or some other remote location. Still, others manage the subscriber equipment and billing requirements.

This method of improving the fault-tolerance, that is, by adding redundant hardware at the head-end, has an additional feature that provides for the sharing of the processing loads prior to storing the processed data. Specifically, some of the managing modules have the ability to process data on more than one head-end controller at a time. These managing modules are termed "distributed," since each instance of the managing module is processing a subset of the session-state data. Session-state data that is processed by a distributed managing module is concurrently being processed at the primary head-end controller processor 135-1 and processed at the secondary head-end controller processor 135-2.

The session-state data processed by the primary head-end controller 130-1 is stored on its dedicated fixed storage device 137-1 and memory 136-1. Likewise, the session-state data processed by the secondary head-end controller 130-2 is stored on its dedicated fixed storage device 137-2 and memory 136-2. In this fashion, the time to process data via a distributed managing module is beneficially reduced almost in half.

Another method of processing session-state data is by means of a "non-distributive" managing module. Session-state data processed by non-distributive managing modules is processed at a primary head-end controller where the processor is said to be in an active mode and corresponding to that specific data being processed. The secondary head-end controller in this instance is not processing that specific data concurrently and is set in a standby mode. In this manner, a processor 135 of a head-end controller 130 is always processing data as a primary processor while also serving as a secondary head-end controller in standby mode.

In the event the second head-end controller 130-2 should become inoperative during the processing of data executed by either of these distributive managing modules, then the first head-end controller 130-1 will continue to process the session-state data. This is accomplished by retrieving the latest updates of the previously processed and stored data on the second head-end controller's storage device 137-2. In this first embodiment, some minimal amount of session-state data may be lost. However, by utilizing a centrally

networked storage device as depicted the second embodiment shown in FIGS. 2A and 2B, all previous session-state data is available for continued processing by a secondary head-end controller, in the event of a primary head-end controller failure.

In the instance where the data being processed is through a non-distributive managing module such as the content manager or subscriber manager, then the processing occurs by the head-end controller designated to the access controller associated with the requesting subscriber. For example, a subscriber sends a request for video information that is routed through the first access controller 140-1. The first head-end controller 130-1 is assigned to that access controller 140-1. As such, the primary head-end controller 130-1 and its corresponding subscriber manager located thereon processes such a request and then stores the session-state data produced by the subscriber manager on the corresponding storage device 137-1 and memory device 136-1.

At the same time, the subscriber manager located on the second head-end controller 130-2 does not process the data but remains in a standby mode. If the primary head-end controller 130-1 becomes inoperative while processing such data, then the access controller 140 dedicated to the requesting subscriber diverts data traffic to the secondary head-end controller 130-2, and the subscriber manager on the secondary head-end controller 130-2 becomes active. The subscriber manager on the secondary head-end controller 130-2 then retrieves the stored session-state data from the primary head-end storage device 137-1 and memory device 136-1 for continued processing. Again, a minimal amount of session-state data may not be recovered in the instance where each head-end controller 130 utilizes a local storage device 137 as depicted in FIG. 1A and 1B. However, all the executed data will be recoverable in the instance where a centrally networked storage device 139 is utilized, as depicted in FIG. 2A and 2B.

Common to all of these managing functions is that data at each head-end controller 130 is constantly being processed by it's respective processor 135 and stored either temporarily on it's random access memory 136 or permanently on it's permanent storage device 137 or 139 of FIGS 1A or 2A. Therefore, in the event one of the head-end controllers 130 becomes inoperable during a subscriber session, the redundant head-end controller 130 can access the memory and storage device of the inoperable head-end

controller. Where the storage device 137 is local to each head-end controller 130, processing the session-state data will continue with minimal interruption to the system or subscriber. In the instance where the storage device 139 is centrally networked to each head-end controller 130, then processing of session-state data will resume without any
5 interruption.

FIGS. 3A and 3B (collectively FIG. 3) depict a flowchart of a method 300 that facilitates fault tolerance at a head-end of an interactive information distribution system. A normal mode of operation, i.e., no failure occurring at the primary head-end controller, is depicted by the flowchart of FIG. 3.

10 A subscriber in an interactive information distribution system, such as a video-on-demand system, starts the method at step 301, and the method proceeds to step 302 where the subscriber requests video information from their subscriber equipment, i.e., a set-top box. In step 304, the set-top box sends the request through the back channel, to the access system controller dedicated to such set-top box. In step 306, the access system controller
15 receives the subscriber's request and then signals the primary head-end controller to initiate a subscriber session. In step 308, the primary head-end controller initiates the subscriber session through the managing modules that regulate the stream server's video transmissions at the head-end.

In step 310, non-distributed managing modules of the primary head-end controller
20 and the distributed managing modules of both the primary and secondary head-end controllers process the session-state data. In step 312, the session-state data is stored on the primary head-end memory and storage devices. The memory and storage devices include RAM and either a local or centrally networked disk drive, as depicted in FIGS. 1A and 2A, respectively. The session-state data stored on the primary head-end storage devices is
25 replicated, in step 314, onto the storage devices of the secondary head-end controller, where the storage devices are local to each head-end controller, as depicted in FIGS. 1A and 1B. The replication is performed in frequent intervals to provide current session-state data for redundancy purposes in the event of a primary head-end failure. Where the storage device is centrally networked between the primary and secondary head-end controllers as in FIGS.
30 2A and 2B, then the step 314 is excluded from the routine 300. Step 314 is excluded in this

instance since the secondary head-end controller has ubiquitous access to the central storage device that concurrently stores the session-state data created by all the head-end controllers, including the primary head-end controller.

During the processing of the session through the managing modules, in step 316, the plurality of managing modules direct the stream server to retrieve the requested video information from the video content repository and stream the video information to the designated access controller. In step 318, the access controller receives the streamed video information from the server, and then modulates the video information for transmission through the forward channel to the set-top box of the subscriber. The primary head-end controller continues to manage the subscriber's session at step 320, until the session ends, in step 322.

FIG. 4 depicts a flowchart of a method for continuing a subscriber session during a failure mode of operation at a head-end controller. In a failure mode of operation, a failure occurs at the primary head-end controller designated to facilitate a subscriber session. In this instance, the secondary head-end controller retrieves the processed session-state data from the fixed storage device and memory of the primary head-end controller in order to continue the subscriber session.

The failure mode of operation begins at step 401 and proceeds to step 402, where a failure occurs at the primary head-end controller. The failure is detected in an interrupt driven manner from the dedicated access controller previously interfacing with the primary head-end controller. In step 404, the access controller contacts the secondary head-end controller to continue the subscriber's video session. The routine 400 proceeds to step 406 where the managing modules of the secondary head-end controller retrieve the latest session-state data available from the primary head-end controller. The data is stored in either the memory and storage devices of the primary head-end controller shown in FIGS. 1A and 1B, or the centrally networked storage device shown in FIGS. 2 and 2B.

Once the session-state data from the primary head-end controller is accessed, in step 408, both the non-distributed and distributed managing modules of the secondary head-end process the recovered session-state data. In this manner, there is minimal interruption to the subscriber's video session. Thereafter, in step 410, the secondary head-end controller

continues the subscriber's session and stores the session-state data on the memory and storage devices of the secondary head-end controller.

During the execution of the session by the managing modules, in step 412, the managing modules continue to direct the stream server to retrieve the requested video
5 information from the video content repository, and stream the video information to the access controller designated to the subscriber. In step 414, the access controller receives the streamed video information from the server and modulates the video information for transmission through the forward information channel, to the set-top box of the subscriber. The secondary head-end controller continues to manage the subscriber's session in step
10 416, until the session ends, at step 418.

Thus, fault tolerance is achieved by adding additional head-end controllers to each head-end. This allows for the replicating of processed information onto the memory and storage devices of the secondary head-end controller. By utilizing the methods and apparatus disclosed herein, a subscriber using this improved fault tolerant system will be
15 ensured minimal disruptions to their viewing requirements. Furthermore, the implementation of load sharing between two or more head-end controllers decreases the processing time, increases the availability of video information and allows for greater expansion of viewership. Although various embodiments that incorporate the teachings of the present invention have been shown and described in detail herein, those skilled in the
20 art can readily devise many other varied embodiments that still incorporate these teachings.

CLAIMS

What is claimed is:

1. A method of distributing and sharing processing loads and increasing fault tolerance between provider equipment and subscriber equipment of an interactive information
5 distribution system, comprising the steps of:
 requesting video information from said subscriber equipment;
 executing a video session from at least one managing module on a primary head-end controller;
 dedicating at least one secondary head-end controller having said at least one
10 managing module as a reserve resource for executing said video session;
 storing session-state data from said executed video session on at least one storage device; and
 streaming said video information to said requesting subscriber equipment during a normal mode of operation.
15
2. The method of claim 1, wherein a function of said at least one managing module is distributed further comprises the step of:
 executing said video session from said at least one managing module on said
primary head-end controller and said at least one secondary head-end controller.
20
3. The method of claim 2, wherein a function of said managing module is not distributed amongst said primary and secondary head-end controller and said method comprises the step of:
 executing said video session from said at least one managing module on said
25 primary head-end.
4. The method of claim 3, comprising the steps of:
 processing said session-state data through said at least one distributed managing
module concurrently on said primary head-end controller and said at least one
30 secondary head-end controller, wherein said at least one distributed managing module

on said primary head-end controller and said at least one secondary head-end controller is in an active mode; and

- processing said session-state data from said at least one non-distributed managing module on said primary head-end controller, wherein said at least one non-distributed managing module on said primary head-end controller is in an active mode, and
5 wherein said at least one non-distributed managing module on said secondary head-end controller is in a standby mode.

5. The method of claim 4, a method comprising the steps of:

- 10 processing said session-state data produced by said primary head-end controller via said at least one secondary head-end controller in a failure mode of operation, wherein said primary head-end controller becomes inoperative.

6. The method of claim 5, comprising the steps of:

- 15 streaming video information from a stream server to an access controller in said normal mode of operation, wherein said primary head-end controller manages said video session between said stream server and said plurality of access controller; and
streaming video information from said stream server to said access controller in said failure mode of operation, wherein said secondary head-end controller manages said
20 video session between said stream server and said access controller.

7. The method of claim 1, comprising the steps of:

- storing said session-state data produced by said primary head-end controller on said at least one storage device coupled to said primary head-end controller; and
25 storing said session-state data produced by said at least one secondary head-end controller on said at least one storage device coupled to said primary head-end controller.

8. The method of claim 7, comprising the step of:

replicating said stored session-state data from one of said plurality of storage devices coupled to said primary head-end controller, to each of the remaining storage devices of said plurality of storage devices coupled to said at least one secondary head-end controller; and

- 5 wherein said at least one secondary head-end controller retrieves said session-state data executed by said managing modules of said primary head-end controller for continuing said video session with said subscriber equipment.

9. The method of claim 1, further comprising the steps of:

- 10 storing said session-state data produced by said primary head-end controller on a memory device coupled to said primary head-end controller; and

 storing said session-state data produced by said at least one secondary head-end controller on said memory device coupled to said primary head-end controller.

- 15 10. The method of claim 9, comprising the step of:

 replicating said stored session-state data from said memory device coupled to said primary head-end controller, to at least one memory device coupled to said at least one secondary head-end controller; and

- wherein said at least one secondary head-end controller retrieves said session-state data executed by said managing modules of said primary head-end controller for
20 continuing said video session with said subscriber equipment.

11. Apparatus for distributing information through a network comprising:

- 25 a stream server;

 a plurality of head-end controllers, coupled to said stream server, for managing a video session; and

 a plurality of access controllers, coupled to said plurality of head-end controllers, for interacting with said subscriber equipment during said video session to responsively

provide video information to a network upon a request for video information from the network.

12. The apparatus of claim 11, wherein each head-end controller of said plurality of head-
5 end controllers comprises:
a plurality of managing modules for executing said video session;
a processor for processing session-state data produced by said plurality of managing
modules; and
memory devices, coupled to said processor, for temporarily storing said session-
10 state data.
13. The apparatus of claim 12 wherein said plurality of head-end controllers comprises a
primary head-end controller and at least one secondary head-end controller.
- 15 14. The apparatus of claim 13, wherein:
in a normal mode of operation, said primary head-end controller interacts with said
stream server to generate said video information, and said at least one secondary head-
end controller remains in a standby mode; and
in a failure mode of operation, said primary head-end controller is inoperative, and
20 said at least one secondary head-end controller interacts with said stream server to
produce video information.
15. The apparatus of claim 14, wherein said plurality of managing modules comprise:
at least one distributed managing module, for processing session-state data through
25 both primary head-end controller and said at least one secondary head-end controller
concurrently; and
at least one non-distributed managing module, for processing session-state data by
said primary head-end controller.
- 30 16. The apparatus of claim 15, wherein:

in a failure mode of operation, a portion of said plurality of access controllers coupled to said inoperable primary head-end controller interface with said secondary head-end controller, whereby all of said plurality of access controllers are interfacing with said at least one secondary head-end controller, to responsively interact with said
5 network.

17. The apparatus of claim 16, wherein:

in a failure mode of operation, said at least one distributed managing module and said at least one non-distributed managing module executes said video session through
10 said at least one secondary head-end controller.

18. The apparatus of claim 17 further comprising:

a centrally networked storage device coupled to said primary head-end controller and said at least one secondary head-end controller, for centrally storing said session-state data produced by said plurality of managing modules; and
15

in said failure mode of operation, said at least one secondary head-end controller retrieves said session-state data stored on said centrally networked storage device by said primary head-end controller, for continued interaction with said stream server to provide said video information to said network.
20

19. The apparatus of claim 17, further comprising:

a plurality of local storage devices, coupled to said primary head-end controller and said at least one secondary head-end controller, for locally storing said session-state data produced by said plurality of managing modules.
25

20. The apparatus of claim 19, wherein:

said session-state data is replicated from one of said plurality of local storage devices coupled to said primary head-end controller, and stored on the remaining plurality of local storage devices of said at least one secondary head-end controller.
30

21. The apparatus of claim 120, wherein:

in a failure mode of operation, said at least one secondary head-end controller
retrieves said replicated session-state data stored on said remaining plurality of storage
devices, for continued interaction with said stream server to provide said video
5 information to said network.

ABSTRACT

A method and apparatus for load sharing and increasing fault tolerance at a head-end of an interactive information distribution system during a subscriber's video session. A
5 primary head-end controller processes video session data to manage the distribution of
streamed video information to a subscriber. In the event of a failure at the primary head-end
controller, a secondary head-end controller continues to process session-state data, with
minimal interruption to the subscriber's video session.

Copyright © 1999 by AT&T Knowledge Ventures, L.P.

FIG. 1A


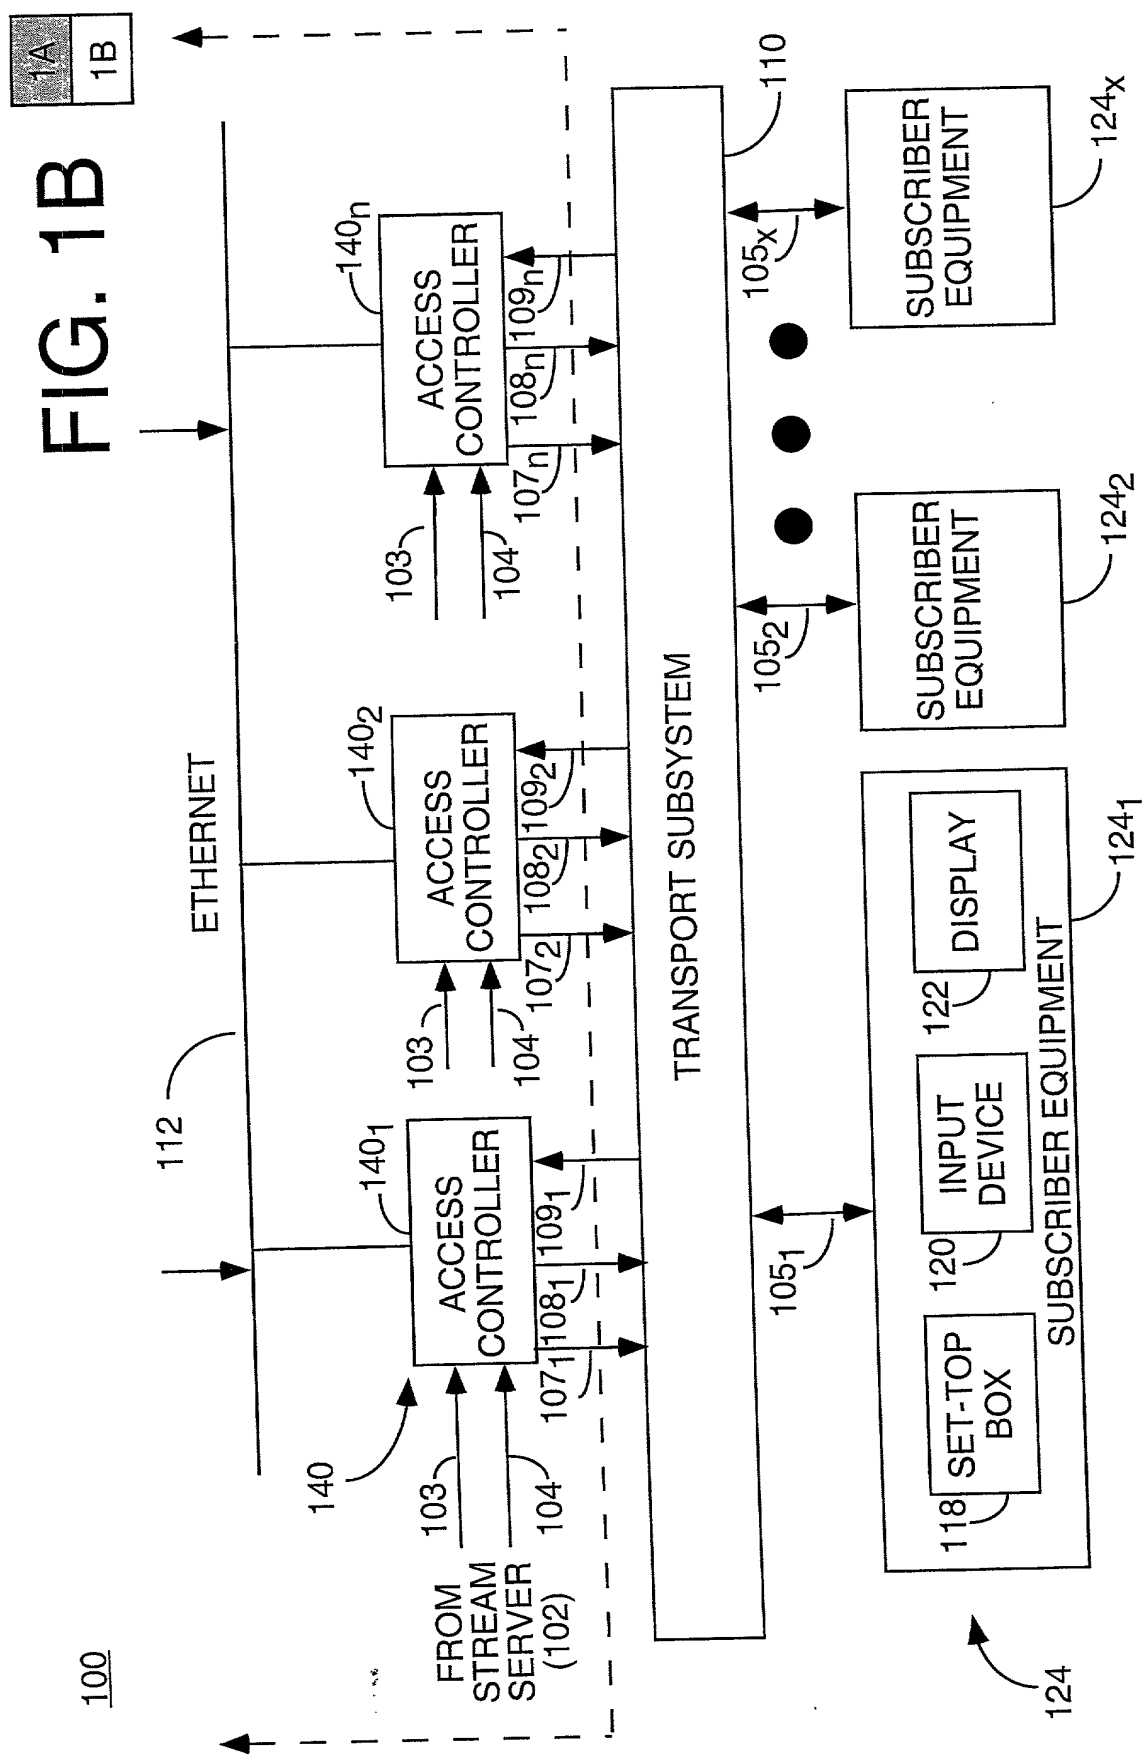


FIG. 1B



100

1A
1B

FIG. 2A

100





FIG. 2B



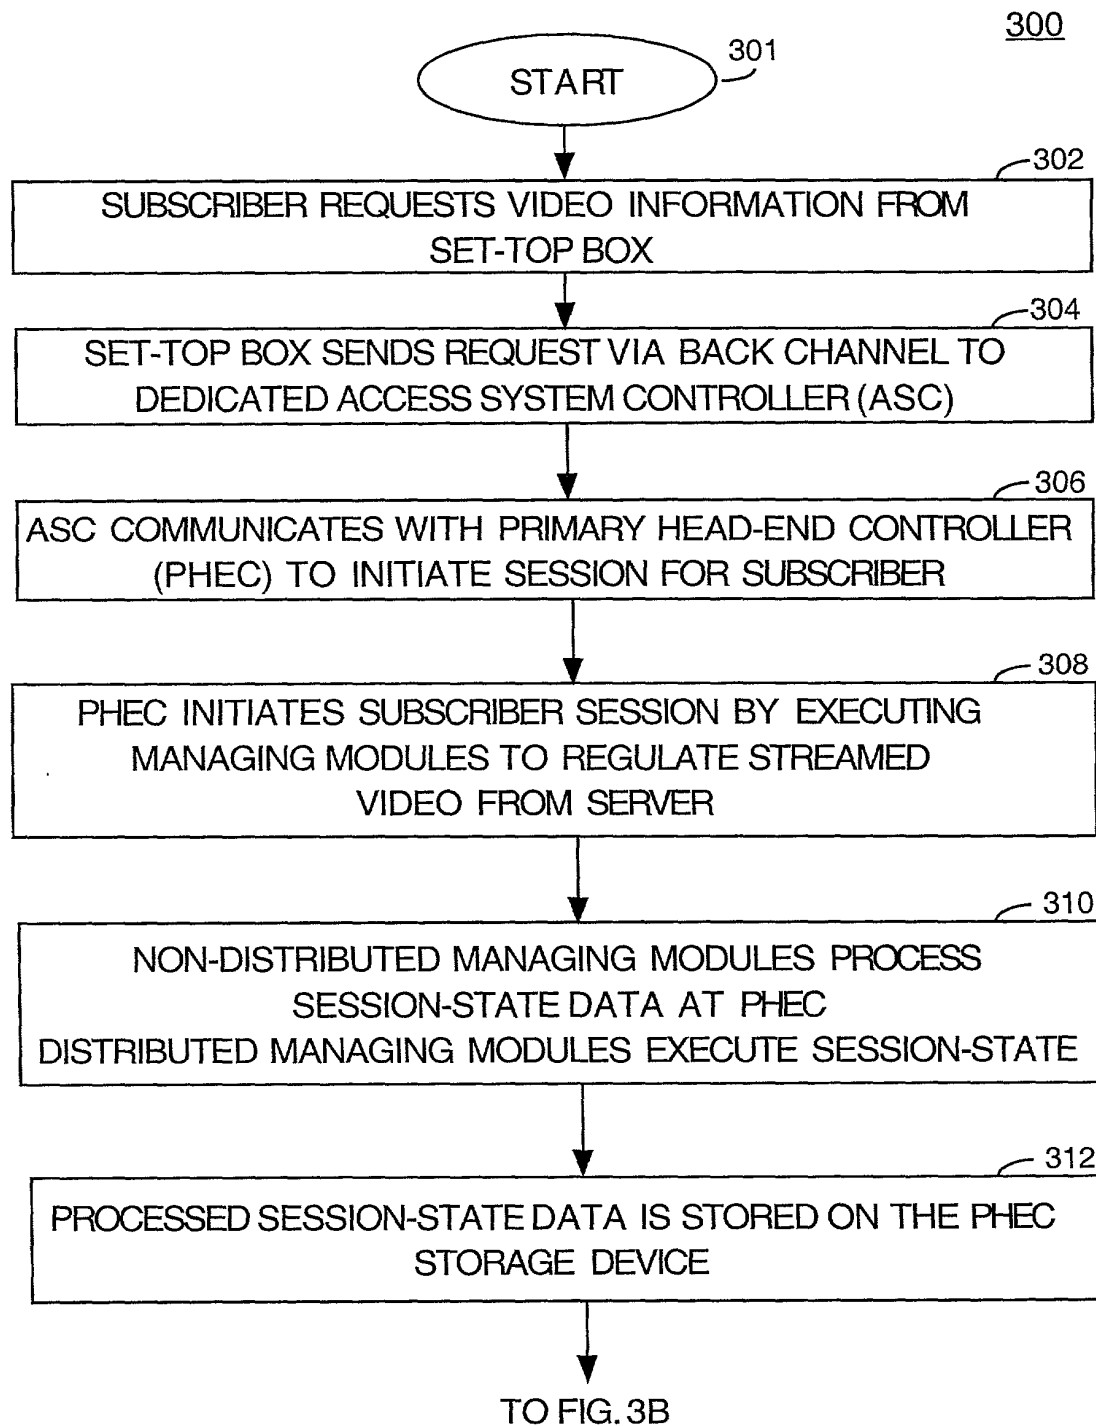



FIG 3A



FROM FIG. 3A


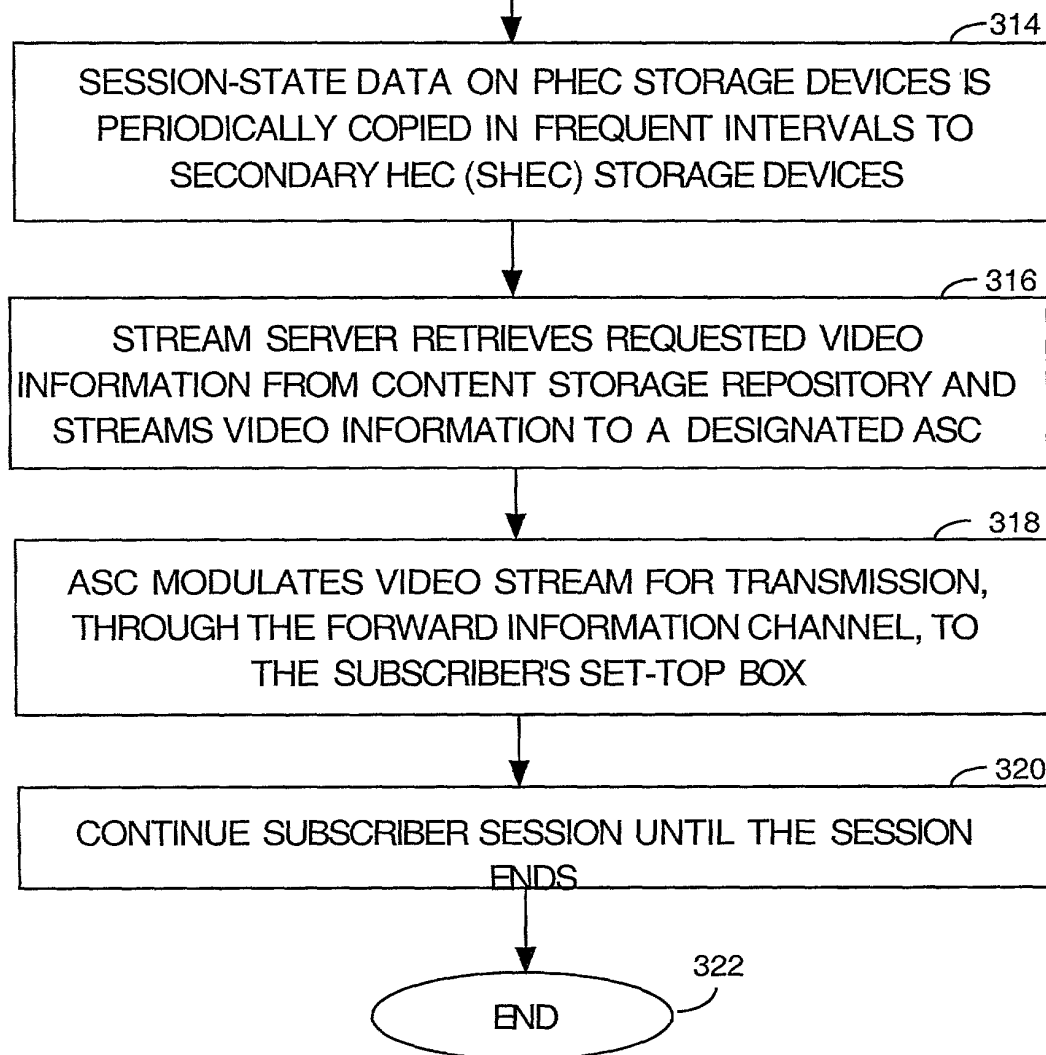


FIG 3B

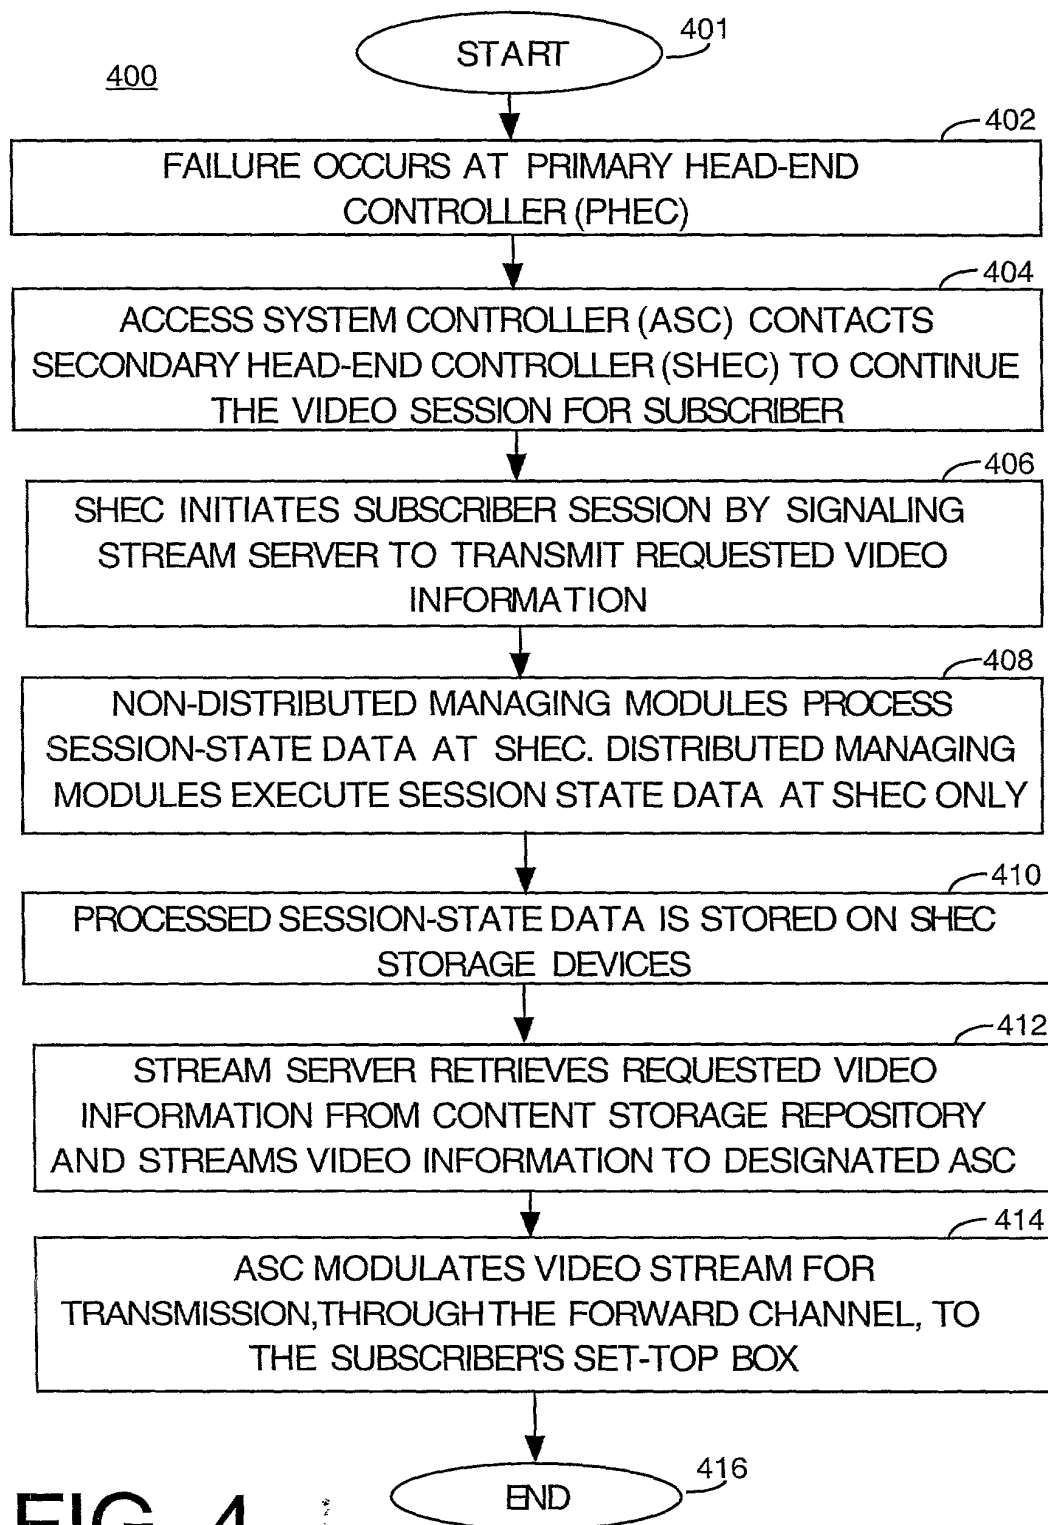


FIG. 4

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I hereby declare that:

This declaration is of the following type:

- ☒ original
- ☐ divisional
- ☐ continuation
- ☐ continuation-in-part

INVENTORSHIP IDENTIFICATION

My residence, post office address and citizenship are as stated below next to my name. I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

TITLE OF INVENTION

METHOD AND APPARATUS OF LOAD SHARING AND FAULT TOLERANCE IN AN
INTERACTIVE VIDEO DISTRIBUTION SYSTEM

SPECIFICATION IDENTIFICATION

The specification of which:

- ☒ is attached hereto
- ☐ was filed on _____, under Serial No. _____, executed on even date herewith; or
- ☐ Express Mail No. _____ (as Serial No. not yet known)
- ☐ and was amended on _____ (if applicable)
- ☐ was described and claimed in PCT International Application No. _____ filed on _____ and as amended under PCT Article 19 on _____.

ACKNOWLEDGEMENT OF REVIEW OF PAPERS AND DUTY OF CANDOR

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose all information I know to be material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56, and which is material to the examination of this application; namely, information where there is a substantial likelihood that a reasonable Examiner would consider it important in deciding whether to allow the application to issue as a patent, and

- ☐ In compliance with this duty, there is attached an Information Disclosure Statement in accordance with 37 CFR §1.98.

PRIORITY CLAIM (35 U.S.C. §119)

I hereby claim priority benefits under Title 35, United States Code, §119, of any foreign application(s) for patent or inventor's certificate or of any PCT international application(s) designating at least one country other than the United States of America or of any United States Provisional Application(s) listed below, and have also identified below any foreign application(s) for patent or inventor's certificate or any PCT international application(s) designating at least one

country other than the United States of America filed by me on the same subject matter having a filing date before that of the application(s) of which priority is claimed.

☐ No such applications have been filed.

☒ Such applications have been filed as follows:

A. Prior foreign/PCT/provisional application(s) filed within 12 mos. (6 mos. for design) prior to this application, and any priority claims under 35 U.S.C. §119

<u>Country/PCT</u>	<u>Application No</u>	<u>Date Filed</u>	<u>Priority Claimed</u>
U.S.A.	Provisional Application No. 60/127,337	April 1, 1999	<input checked="" type="checkbox"/> YES <input type="checkbox"/> NO <input type="checkbox"/> YES <input type="checkbox"/> NO <input type="checkbox"/> YES <input type="checkbox"/> NO

B. All foreign application(s), if any, filed more than 12 mos. (6 mos for design) prior to this U.S. application

Country:

Application No:

Filing date:

**PRIOR U.S. APPLICATION(S) FOR WHICH BENEFIT
UNDER 35 U.S.C. §120 IS CLAIMED**

<u>Serial No.</u>	<u>Filing Date</u>	<u>Status</u>		
		<u>Patented</u>	<u>Pending</u>	<u>Abandoned</u>
		<input type="checkbox"/> patented	<input type="checkbox"/> pending	<input type="checkbox"/> abandoned
		<input type="checkbox"/> patented	<input type="checkbox"/> pending	<input type="checkbox"/> abandoned
		<input type="checkbox"/> patented	<input type="checkbox"/> pending	<input type="checkbox"/> abandoned

POWER OF ATTORNEY

I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Charles Lee Thomason Reg. No. 31,431
Raymond R. Moser Jr. Reg. No. 34,682
Joseph Pagnotta Reg. No. 39,322
Kin-Wah Tong Reg. No. 39,400

Eamon J. Wall Reg. No. 39,414
R. Gale Rhodes Jr. Reg. No. 19833
Wan Yee Cheung Reg. No. 42,410

Send correspondence to:

Thomason, Moser & Patterson LLP
The Galleria -- 2d Floor
2-40 Bridge Ave.
Post Office Box 8160
Red Bank, NJ 07701

Direct telephone calls to:

Raymond R. Moser, Jr.
732-530-9404

DECLARATION

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and, further, that these

statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Sec. 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patents issued thereon.

Full name of **sole or first** inventor: Tom Giammarresi

Inventor's signature: Tom Giammarresi Date: 12/6/99
Residence: 30567 Del Valle Place, Union City, CA 94587
Post Office Address: same as above

Country of Citizenship: USA

Full name of **second** inventor:

Inventor's signature: _____ Date: _____
Residence: _____
Post Office Address: _____

Country of Citizenship: _____

Full name of **third** inventor:

Inventor's signature: _____ Date: _____
Residence: _____
Post Office Address: _____

Country of Citizenship: _____

Full name of **fourth** inventor:

Inventor's signature: _____ Date: _____
Residence: _____
Post Office Address: _____

Country of Citizenship: _____

Full name of **fifth** inventor:

Inventor's signature: _____ Date: _____
Residence: _____
Post Office Address: _____

Country of Citizenship: _____

Full name of **sixth** inventor:

Inventor's signature: _____ Date: _____
Residence: _____
Post Office Address: _____

Country of Citizenship: _____